arXiv:1703.07595v1 [cs.CV] 22 Mar 2017

Can you tell where in India I am from? Comparing humans and computers on fine-grained race face classification.

Harish Katti, S. P. Arun
Centre for Neuroscience
Indian Institute of Science
harish2006@gmail.com, sparun@cns.iisc.ernet.in

Abstract

Faces form the basis for a rich variety of judgments in humans, yet the underlying features remain poorly understood. Replicating these judgments using computers can enable a vast variety of practical applications. While coarse categories such as race and gender are relatively easier to deduce and are widely studied, they do not constrain the possible features used by humans since these classes differ in a large number of features. In contrast, fine-grained distinctions within a race might more strongly constrain the possible features, but are relatively less studied. Fine-grained race classification is also interesting because even humans may not be perfectly accurate on these tasks. This offers a unique opportunity to compare errors made by humans and machines, in contrast to standard object detection tasks where human performance is nearly perfect. Here we set out to address this gap by developing, benchmarking on humans and machines and publicly releasing novel Centre for Neuroscience Indian face dataset (CNSIFD) face database of close to 1650 diverse Indian faces labeled for fine-grained race (South vs North India) as well as for age, weight, height and gender. We then performed an extensive behavioral study on close to 130 human subjects who were asked to categorize each face as belonging to Northern or Southern state in India. We then compared human performance on this task with that of computational models trained on the ground-truth labels.

Our main results are as follows: (1) Humans show extremely consistent performance across faces (average accuracy: 63.6%), with some faces being consistently classified with >90% accuracy and others consistently misclassified with <30% accuracy; (2) Models trained on ground-truth labels showed slightly worse performance (average accuracy: 62%) but showed higher accuracy (72.2%) on faces classified with >80% accuracy by humans. This was true for models trained on simple spatial and intensity measurements extracted from faces as well as deep neural networks trained on race or gender classification; (3) Using over-complete banks of features derived from each face part, we found that mouth shape was the single largest contributor towards fine-grained race classification, whereas distances between face parts was the strongest predictor of gender- Our results show that computational models can explain a variety of face attributes and we hope that this study spurs further work on finer-grained classification of faces.

1. Introduction

> "*Just realize where you come from: this is the essence of wisdom.*"
> —*Tao Te Ching, v. 14*

We make a rich variety of judgments on faces ranging from gender, race, personality, emotional state etc. Yet the features that enable this range of judgments remain poorly understood but are of great practical significance since they can enable a variety of applications. Among these judgments is race, which has typically been studied at the coarse level (Caucasian/Black/Asian) [1][14]. While this is relatively easy since many features can vary between coarse race categories, finer race categories provide stronger constraints on the features used by humans since fewer features vary across these classes. This problem has received relatively little attention. The few studies that exist deal with Asian faces [16][15] [2] and there is very little work on Indian faces.

People of Indian origin account for over 12% of the worlds population and have a remarkable diversity of cultural, social and linguistic systems. Just as the Indian geography can be divided roughly into northern and southern parts, this division is also accompanied by stereotyped facial appearance. Our operational definition of north and south Indian regions is shown in Figure 2. The ethnicity of Indians is often a reflection of distinct cultural, social and linguistic traditions and hence it is socially advantageous to know a person's race. This problem is non-trivial due to the

large variability in visual appearances of people from the same geographical region as illustrated in Figure 1.

Discriminating between north and south Indians remains unexplored in the compter vision domain and lacks proper operational definitions as well. Coarse or fine grained race categorisation is typically studied using discrete labels that do not reflect the underlying continuum of face features, nor do they reflect the graded changes in human perception for these faces. Some recent studies on race/ethnicity classification have modeled attributes such as cosmetics, jewelry and even hair style [16], though such attributes could be informative, they can be easily manipulated or obscured during image acquisition [4]. We emphasize on the core problem of race classification using variations in intensity and spatial measurement information available on faces.These facial attributes remain relatively invariant to personal choices and cultural changes brought about by globalisation for example. We have three main contributions in this paper.

First we introduce the problem of discriminating between North and South Indian faces and show that human participants can do this task well above chance. For this purpose, we performed behavioural experiments where close to 130 human subjects categorised these faces in speeded north vs south categorisation tasks. Second, we show that race labels as self-declared by an individual and can be qualitatively different from the fine-grained race perceived consistently by human observers. Thirdly, we have evaluated the performance of many relevant computational models and compared them with that of humans, to get insights into plausible feature extraction schemes used by humans to perform race discrimination. To the best of our knowledge, our dataset not only captures the diversity of Indian faces more faithfully, but also is very rich as we gathered multiple attributes such as self-declared race, gender, age, height and weight of individuals. We will make our dataset, feature extraction schemes and implementation code available in the public domain to spur further research in this topic.

1.1. Related work

We first review literature from experimental studies in human vision that have explored various facets of race discrimination from face images. Coarse race classification between groups like caucasians and black have been explored in psychology [1][4] and interestingly it has been shown that humans can reliably perform this task in the absence of salient cues such as skin colour, expressions, cosmetics, ornaments or attributes such as hair style [1]. This again points to the fact that the core problem race classification from face images is independent of these attributes.

Automated methods for race classification have shown upto 70-80% classification performance in some race classification problems [4]. Although interesting datasets do exist for Indian faces ,they focus on identity and age discrimination [13][12] or on celebrity recognition [11]. Models trained with feature extraction schemes using Local binary patterns, wavelets and Gabor filter banks have been trained to near human level performance [4]. More recently deep convolutional networks (CNN) [6] have given impressive performance in analogous tasks such as discriminating between Chinese, Korean and Japanese faces [16] and features derived from CNNs have largely replaced other feature encoding schemes. There are several unexplored aspects in such studies, for example expert annotated ground truth is discrete whereas we show here that perceived race is really a continum like many other face attributes and for algorithms to perform in a human like manner, they should also be able to capture variability exhibited by human participants in race discrimination. Furthermore, features such as Local binary patterns [9], responses of gabor filterbanks [3] and wavelets, do not lend for easy interpretation and there is often a disconnect between models that perform well and those that are interpretable. We address these concerns in our paper.

1.2. Overview

In the following sections we first introduce our face dataset for North Vs South Indian classification in 2. We then describe experiments where participants performed speeded race discrimination on our face dataset and summarise results from this experiment in 3. In section 4 we describe various feature extraction schemes and models for automatic race categorisation and evaluate the performance of models and compare it with human performance in section 5. We then summarise our study in section 6.

2. Dataset

We first gave an operational definition for North Indians as people who belong to states highlighted in black in Figure 2, likewise we define South Indians as people who belong to states highlighted in green in the same figure. Our face dataset has a total of 1647 Indian faces of which 459 faces have been collected from volunteers who gave informed consent and agreed that their faces could be used for research. These volunteers were photographed in high resolution declared that they as well as both their parents belong to either a north Indian or south Indian state. The remaining faces were obtained from the internet by searching for face images on the web using combinations of typical first names (north: *eg; birender, payal*, south: *eg; jayamma, thendral*) and last names like (north: *eg; khushwaha, yadav*, south: *eg; reddy, iyer*). We obtained these names by collating lists of names reported by upto six volunteers and retaining only the most frequently mentioned typical first and last names for each state and we made use of Google Image search API's to search for combinations of these first

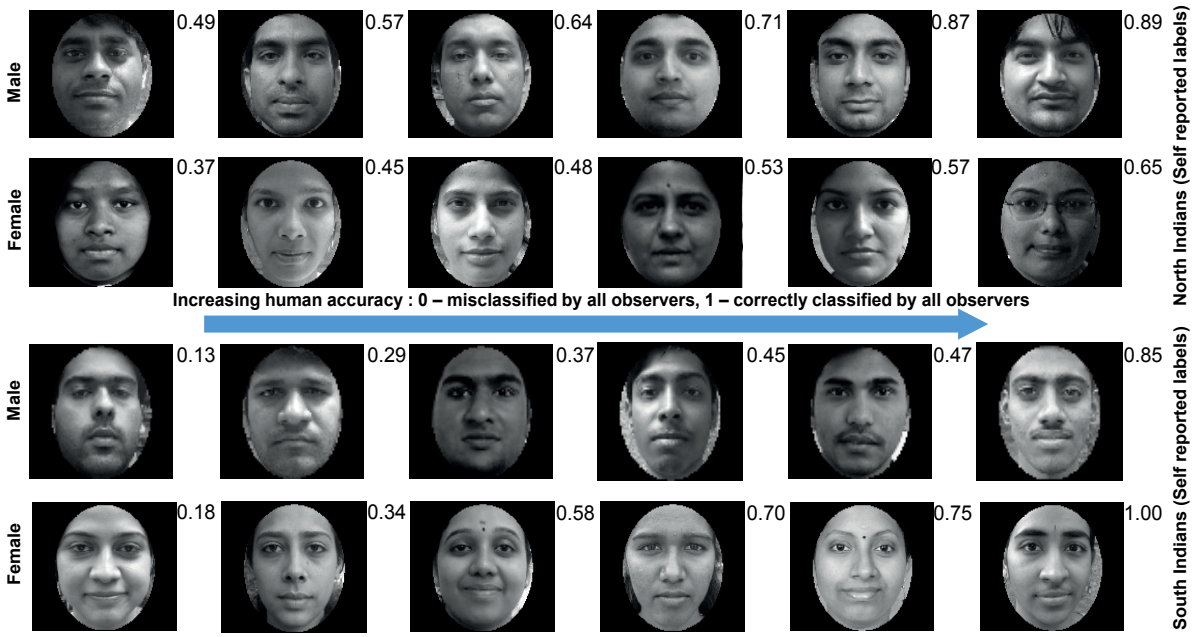

Figure 1: Illustrative north and south Indian faces from our dataset. The mean categorisation accuracy of human observers is shown above each face. Faces are segregated into north/south based on self-declared labels from participants. Faces wth average accuracy close to 1 indicate that the race attribute label matched closely with self-declared labels. Figure best seen in high-resolution in the digital version.

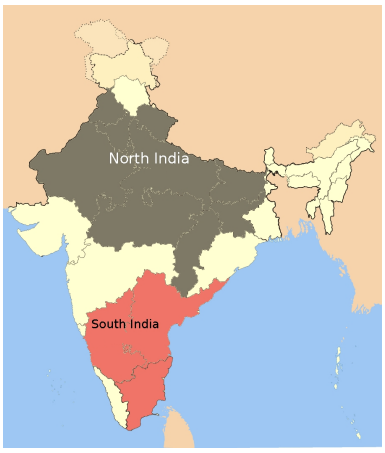

Figure 2: Left : Operational definition of North (black) and South (red) regions marked on the Indian map. We included states that are generally agreed to be part of Northern or Southern parts of India, and excluded states that have unique or distinctive cultural identities (e.g. Kerala, West Bengal, Assam). This is merely an operational definition for the purposes of the study, but the fact that independent sets of subjects were able to categorize faces with high accuracy confirms the validity of our definition.

and last names. For faces crawled from the web, frontal faces were detected using the CART face detector provided in Matlab's computer vision toolbox and faces in high resolution (>150 x 150 pixels) and for which atleast three out of four annotators had an agreement on the race label were selected and retained for further experiments. These faces were then additionally annotated for gender labels as well. For the 459 images contributed by volunteers, we additionally instructed them to report demographic information such as age, height and weight if possible. We then normalised all the faces by registering 76 facial landmarks [8] and then rotating and scaling faces so that the mid-point between eyes on all faces co-incide and so that the vertical distance from chin to eye-brow is 250 pixels. These details are summarised in Table 1. Upto eight annotators then inspected faces crawled from the web and assigned them a north indian or south indian label, we chose faces that had a majority consensus. We verified post-hoc that the part of our dataset containing 1188 faces with consensus based region labels were nearly as hard (63.88%) to categorise into north and south as the subset of 459 faces (62.3%) for which volunteers had declared their race.

3. Behavioural experiment

We setup a speeded categorisation task where we showed subsets of faces from our dataset to Indian participants and

Face set	Total	M	F	N	S	Other
Consent	459	260	199	140	209	110
Web	1188	710	478	636	552	0

Table 1: Breakdown of faces in our dataset. The first row **Consent** gives details of faces where volunteers gave informed consent prior to being photographed, many of these faces have other attribute information such as age, height and weight in addition to region and gender labels. The second **Web** row gives details of faces that were gathered from the internet and curated by 2-8 observers who did not participate in subsequent experiments. **M**-Male, **F**-Female, **N**-North, **S**-South

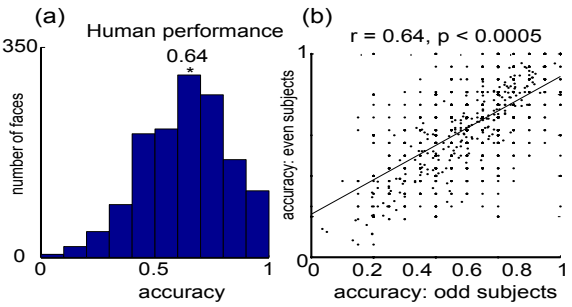

Figure 3: Distribution of human accuracy on north/south categorisation. Accuracy is defined as the fraction of participants who correctly guessed the declared label.

instructed them to indicate on every trial if a face was north or south Indian. All faces were shown in grayscale. Participants were instructed to be fast and accurate and no feedback was given to participants about their performance on a trial-by-trial basis. In all we obtained responses from close to 130 participants for 1423 faces and in this manner we obtained over 16 unique responses for each face. The average human accuracy on correctly categorising faces as north or south Indian was 63.6% and there is considerable variability in human accuracy as shown in Figure 3 (a). This indicates that although humans are well above chance ($p < 0.0005$, one sample t-test against a distribution with mean of 50%), it is still a challenging task. This difficulty arises from the diversity present within male as well as female faces belonging to either north Indian regions or south Indian regions and can be seen in Figure 1. Despite the variability in human accuracy, subjects were nonetheless very consistent in their responses, as evidenced by a high and statistically significant correlation between the average accuracy of two halves of the subjects across all faces ($r = 0.64$, $p < 0.0005$). The high degree of consistency is also illustrated in Figure 3 (b) where we have plotted the average human accuracy from even numbered subjects, against the average accuracy of odd numbered subjects.

4. Models

Brightness of different face parts are informative of the 3D structure of faces and like wise spatial measurements capture essential aspects of overall face shape such as width of face and relative positions of important face parts. Likewise, stereotyped local shape of face parts can hold information about various face attributes. It is also possible that overall skin-tone could be informative as well although its importance in race discrimination is still a matter of debate [1] and this can be captured by face-wide statistics of intensity distribution. Although many visual features have been explored for face attribute discrimination, deep convolutional networks [10][7] have largely outperformed most conventional feature schemes. We have modeled and evaluated all these feature schemes describe them in further detail in this section.

4.1. Spatial and Intensity features

We defined and used two schemes to sample intensity and spatial(metric) measurements on faces, in the first we defined a detailed and interpretable set of face regions as well spatial measurements illustrated in the top row in Figure 4(b). We then registered an active appearance model [8] that had identified the locations of 76 facial landmarks on each face as illustrated in Figure 4 (a). These landmarks were then used to approximate the patches and spatial measurements defined earlier and mean, minimum and maximum intensity values were recorded along with landmark based spatial features as shown in the bottom row in Figure 4(b) and this yielded a set of 23 spatial and 31 intensity measurements. In the second scheme, selected a subset of 26 landmarks such that they still had good coverage over the entire face as well as key parts therein. We then used the Delaunay triangulation method to recover 43 face patches, each of which covered the same region across all subjects. We then extracted the mean, minimum and maximum intensity in each triangular patch as well as 325 all pair-wise distances between the centroids of these triangles. In all this exhaustive method gave 129 intensity and 325 spatial measurements. In addition to the interpretable and exhaustive sampling of intensity and spatial information, we also extract first 6 moments of the image intensity distrubution including mean, variance, skewness and kurtosis. The motivation was to rule out the possibility that models were solving prediction problems merely by picking up some global image statistic and not really processing facial features specifically.

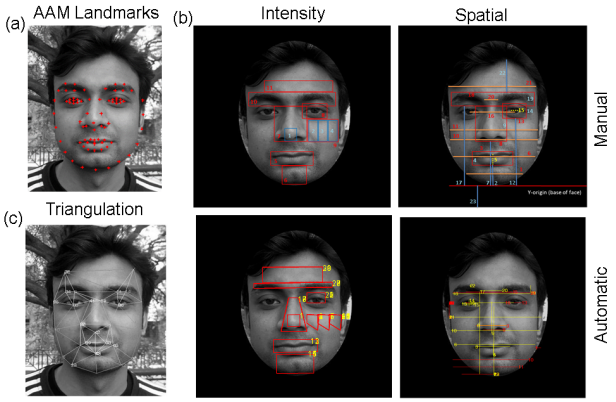

Figure 4: Scheme for sampling spatial and intensity measurements. Panels in the row depict the interpretable spatial and intensity measurements made on each face in the dataset. Bottom row shows the actual measurements made from patches defined on landmarks detected by an active appearance model [8]. The rightmost panel in the bottom row depicts the results of Delaunay triangulation on detected landmarks. These dense patches were then used for the brute force sampling of spatial and intensity measurements. Figure best seen in high-resolution in the digital version.

4.2. Local shape and configural features

We modeled local shape by measuring all pair-wise distances over landmarks detected on left eye, right eye, nose, mouth and face contour for which obtained 36, 36, 66, 153 and 105 such distances respectively. We modeled configural features by taking the centroid-to-centroid distance between face parts including left eye, right eye, nose, mouth, left-contour, right-contour and chin region and for these 7 face locations we obtain 21 pair-wise distance measurements.

4.3. Deep convolutional networks

Since deep convolutional networks (CNNs) have largely replaced previous feature schemes in terms of categorisation performance, we have evaluated the effectiveness of two kinds of CNNs. We have used the VGG-Face [10] face recognition CNN which is has been fine tuned for face recognition, additionally we have also evaluated a CNN trained for age classification [7] and one for gender classification [7] and we refer to these models as CNN-age and CNN-gender hereafter. Both CNN-age and CNN-gender have 3 convolutional layers with 96x7x7, 256x5x5, 384x3x3 filters sizes respectively. This is followed by two 512 node fully connected layers and a single node decision layer. VGG-Face on the other hand is a much deeper network and has 11 convolutional and 3 fully connected layers[10]. Since we consider the penultimate fully connected layer outputs as features for our subsequent analysis, we obtain 512 dimensional feature vectors for each face from CNN-age/CNN-gender and 4096 dimensional feature vectors from VGG-Face.

4.4. Model training and cross validation

We trained linear models that either perform binary classification in the case of discrete race and gender labels, or predict continuous scores as in the case of average human accuracy on north/south race classification, age, height and weight.

Models for binary classification were trained using Linear discriminant analysis implemented in the Matlab® *classify()* function that projects data into a space that maximizes class separability and then mimizes Bayes error by assiging a novel observations x to the class $i \in \{0, 1\}$ that has higher of the two conditional probaiblity $p(i|x)$. Making simplifying assumptions that data comes from multivariate normal distribution and that all class covariance matrices are equal, assignment of a new data observation x to class i is made if,

$$f_i > f_j, i, j \in \{0, 1\}, i \neq j \tag{1}$$

where f_i is given as,

$$f_i = \mu_i C^{-1} x^T - \frac{1}{2} \mu_i C^{-1} \mu_i^T \tag{2}$$

where C is the pooled covariance matrix and μ_i is the mean for class i.

Models for regression were trained using regularised linear regression implemented in the Matlab® *lasso()* function that solves the linear equation,

$$Y = X\beta + \epsilon \tag{3}$$

Where X is an $n \times m$ matrix corresponding to n training faces and m feature dimensions. The solution for coefficient matrix β is obtained by minimizing,

$$\frac{1}{n} |y - X\beta|_2^2 + \lambda |\beta|_1 \tag{4}$$

This term minmizes the difference between predicted and observed values and simultaneously attempts to make the solution for β into a sparse vector. We automated the choicel of the regularisation parameter λ using the inbuilt functionality of Matlab® *lasso()* for this purpose.

To attenuate the effect of noise, we first projected all feature values into Principal component space and retained projections on dimensions that explain 95% of the variance in the data.

Since feature types varied in the number of dimensions, we employed a 10-fold cross validation scheme while training all regression and classification models to avoid overfitting. In the followng section we evaluate the performance of various feature schemes in predicting race, gender and other face attributes.

5. Evaluation and results

We first evaluated the performance of all our model performance on predicting categorical race as well as gender labels. Though our main focus is on fine grained race categorisation, the reasons that we additionally train models for gender are because humans are known to be very good at this task. Gender categorisation gives us a good point of reference to evaluate whether our feature representations are indeed informative. Once we evaluated our models for race and gender categorisation, we investigated their effectiveness in predicting average human accuracy. This allowed us to understand the relative importance of features when humans perform this categorisation task. Lastly we trained models to predict other face attributes such as age, height and weight to better understand the role played by different feature types.

5.1. Predicting race and gender ground-truth labels

To investigate the features useful for race and gender labels, we trained a number of computational models to classify these labels using 10-fold cross-validation. The performance of these models is summarized in Table 2. To determine the best model for race, we looked for the model that gave the highest cross-validated accuracy. Two models yielded equivalent accuracy for race: these were either based on intensity features or CNN-F features (Table 2). To compare the best models with each remaining model, we calculated the number of cross-validated splits in which the best model accuracy was larger than that particular model. Using this method, we found that the best model was better than all other models more than 95% of the time, which corresponds to the standard statistical significance of $\alpha = 0.05$. Using a similar approach we determined the CNN-F features to be the best model for predicting gender labels (Table 2). We note however that the best model for gender had a higher accuracy (85%) compared to the best model for race (62%). This was in fact true for all models as well. This reflects the difference in intrinsic difficulty of the two tasks. To further elucidate the local features that contribute to race & gender labels, we calculated pairwise spatial distances between facial landmark points on each specific part of the face (eye, nose, mouth and contour). This yields an extensive set of measurements for each part that contained a complete representation of its shape. Using this approach we found that mouth shape was the most discriminative part for race classification, whereas inter-part distances were the most informative for gender (Table 2). Interestingly, inter-part distances fared better than exhaustive sets of spatial and intensity features we attribute this to overfitting because the latter models have more degrees of freedom.

To investigate whether low-level brightness statistics are informative about race or gender, we calculated the first five moments of the pixel intensities of each face, and used them as features to predict race and gender labels. Interestingly, while moments were unable to predict race labels, they were slightly above-chance at predicting gender.

Overall we found that CNN models did well in general and VGG-Face in particular was most informative amongst CNN based models, the informativeness of VGG-Face features for face attribute categorisation has been reported in other independent studies such as well[10][5]. All feature types explained relatively less amount of variance in average human accuracy. Some feature types such as VGG-Face perhaps suffered from the curse of dimensionality for some tasks and this led to low performance in predicting average human accuracy (rank = 4).

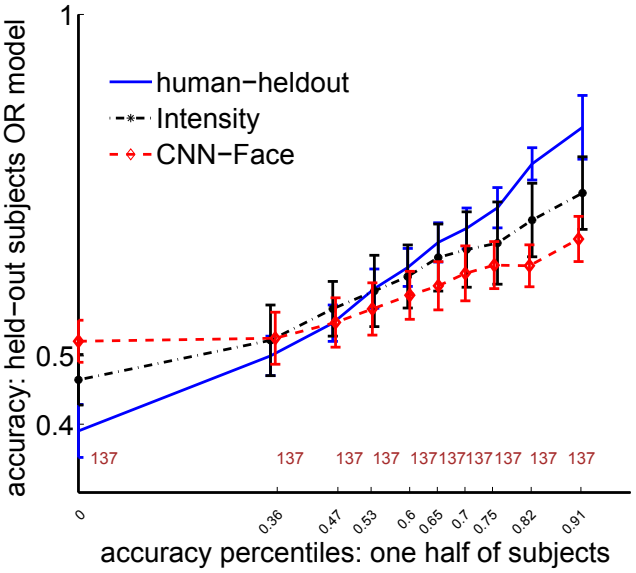

Figure 5: Model and human performance on subsets of faces spanning monotonically increasing equal sized percentile bins of human accuracy of one half of human subjects. The best performing model trained with spatial and intensity features *blue* and those trained with the VGG-Face[10] features *red* are shown along side the average human accuracy of held out half of human subjects. Error bars depict standard deviation about the mean.

5.2. Comparison with humans & modeling human performance

Next we wondered whether faces that were easily classified by humans would be also easy to classify for the models trained on ground-truth labels. To this end, we selected one half of the subjects and used their accuracy on each face to group faces according to their performance. The performance of the other half of the subjects was then used as a model to compare how the accuracy of one group of humans would fare against faces categorized well by another group. To this end, we grouped faces into equal-percentile

bins ranging from faces that were consistently misclassified (accuracy 40%) to those that are highly accurately classified (accuracy = 80%). We then calculated model performance for the two best models for each group of faces. In the resulting plot Figure 5, it can be seen that as faces become easier to classify for one group of humans, they become easier for both other humans as well as the best models. Interestingly, the human data was matched best by the intensity features but not by the CNN-F model. However we note that there is still a 10% gap between human and model performance.

Finally, we trained computational models to predict human performance directly. Since faces classified as North and South could have the same accuracy but be based on different face features, we solved two separate regression problems of the form $y_N = X_N b_N$ and $y_S = X_S b_S$, where y_N and y_S are a vectors of human accuracy on North & South faces, X_N and X_S are matrices of face features corresponding to each face for North & South faces and b_N & b_S are weight vectors representing how the features combine to predict accuracy. The results are summarized in Table3. The best model for human accuracy was one that used spatial and intensity features, but even this model did not reach the split-half reliability for humans. This in turn suggests that humans are using potentially different features for classification.

5.3. Predicting other ground-truth labels (age, weight, height)

Face processing is a ubiquitous task and seldom involves judging race alone. We frequently simultaneously and rapidly evaluate several other important attributes such as gender, age, height and weight. It then becomes plausible that there is a common feature representation that can be flexibly re-weighted to learn decision boundaries for different attributes. To further investigate this, we collected all these additional attributes for as many faces as possible. From amongst the faces that were used for speeded race categorization, 459 faces had age labels, 218 had height, 253 had weight annotations. We find that our baseline scheme involving interpretable spatial and intensity measurements could predict age, weight and height with a significant correlation Table 3. Taken together the performance of these models indicates that overcomplete representations of basic spatial and intensity measurements on faces are highly informative of multiple facial attributes and re-weighting the importance of these features can give appropriate decision boundaries.

5.4. Comparing models across different tasks

We now summarize model performance across diverse face attribute prediction problems and discuss them in further detail. Since there are different numbers of faces with labels/scores for every task and because some models are classifiers and others are regression models, we first compare models within each task and assign them ranks and then compare the rank for every feature type across diverse tasks.

Overall, we have found that models for face attribute prediction when trained with overcomplete representations of spatial and intensity information, perform favorably compared to the best performing CNN derived features. Spatial and intensity feature based models had an average rank of 3.5 over the 6 different prediction problems chosen in this paper and performed favorably against CNN based schemes. Amongst part based local shape features, mouth shape was more informative for predicting average human accuracy (rank=7) than other attributes, nose shape was relatively more informative for age (rank = 6), contour shape is informative for age and height (rank = 3 and 6 respectively) and eye shape was relatively more informative for weight classification (rank = 5) amongst all tasks. We infer this to be task driven reprioritization of local shape information. It is also interesting that the set of 23 spatial and 31 intensity measurements from selective face parts and fiducial landmarks is more informative (rank = 1) for finer grained race categorization as compared to the exhaustive approach of triangulating the entire face and drawing 129 intensity and 325 spatial measurements (ranks = 5). We also find that those features such as mouth shape and intensity measurements for which models did well on north/south ground truth labels (ranks = 8, 2 respectively), also did well on predicting average human accuracy in a rapid categorization task (ranks = 7, 2 respectively). The configuration of face parts as measured by center-to- center distances between parts is highly informative for age prediction (rank = 3). We observed that pretrained CNN models can also exhibit task specificity as in the case models trained with CNN-gender derived features performing much better at gender discrimination (race = 3) than race discrimination (rank = 5).

6. Discussion

Here we have characterized both human and machine performance on a challenging fine-grained race classification task on Indian faces. This is an interesting vision problem because unlike other recognition tasks where human performance is nearly perfect, humans are only 64% correct on this task yet show highly consistent behaviors indicating that faces vary in their intrinsic difficulty of classification. Our dataset also contains race, gender, age, height, and weight that we were able to predict using computational modeling and demonstrate that they are potentially driven by different face features. This problem is therefore interesting and fertile ground for detailed comparisons of machine and human feature representations.

		Race classification: North/South			Gender classification: Male/Female		
Feature	**Dims**	**Df**	**Accuracy**	**Rank**	**Df**	**Accuracy**	**Rank**
#Faces	-	-	1537	-	-	1647	-
Human	-	-	0.64	-	-	~ 1	-
Eye	72	2	0.51(0.005)*	13	5	0.68(0.002)*	11
Nose	66	6	0.53(0.007)*	11	7	0.68(0.002)*	10
Mouth	153	5	0.56(0.004)*	7	6	0.58(0.003)*	15
Contour	105	1	0.51(0.007)*	14	5	0.64(0.003)*	13
ENMC	396	7	0.56(0.005)*	8	16	0.72(0.002)*	9
Inter-part	21	1	0.49(0.008)*	16	5	0.72(0.002)*	8
S	23	4	0.55(0.005)*	9	8	0.68(0.002)*	12
I	31	**9**	**0.62(0.004)**	**2**	12	0.75(0.002)*	7
S+I	54	**12**	**0.62(0.004)**	**1**	19	0.77(0.003)*	6
S+I(Ex)	126	56	0.57(0.007)*	6	55	0.81(0.003)*	2
Moments	7	2	0.50(0.005)*	15	2	0.56(0.001)*	16
CNN-A	512	102	0.60(0.006)*	4	99	0.78(0.004)*	5
CNN-G	512	34	0.59(0.005)*	5	52	0.79(0.002)*	3
CNN-F	4096	**516**	**0.62(0.007)**	**3**	**672**	**0.85(0.004)**	**1**

Table 2: Model performance for race and gender classification. Each feature bank with **Dims** dimensions was projected into PCA space and (Matlab *princomp()*) and only projections on **Df** principal components that explained 95% of the variance in the features were retained. Separate 10-fold cross validated linear classifiers (LDA, Matlab *classify()*) were trained to classify *North/South* and *Gender* labels. Models were trained on 100 repeated splits of 90% of the data and predictions were made on the held out 10%. To find out whether the best model outperformed the rest, we computed the number of times out of the 100 iterations that another model outperformed the best model, all models were trained in each iteration on exactly the same 90% training split. We indicate a * against models whose performance was significantly less than the best model. Average model performance within each task was sorted in descending order to get the model rank **R**. **Df**-degrees of freedom.

	Avg Human Acc			Age			Height			Weight		
Feature	**Df**	**corr**	**R**	**Df**	**corr**	**R**	**Df**	**corr**	**R**	**corr**	**Wt**	**R**
#Faces	-	1647	-	-	459	-	-	218	-	-	253	-
Human	-	0.76	-	-	-	-	-	-	-	-	-	-
Eye	4	0.14(0.006)*	13	4	0.01(0.041)*	13	2	0.24(0.021)*	10	5	0.36(0.013)*	5
Nose	6	0.14(0.005)*	12	5	0.18(0.021)*	6	5	0.32(0.019)*	9	5	0.32(0.013)*	7
Mouth	4	0.17(0.005)*	7	4	0.00(0.029)*	14	5	0.24(0.018)*	11	4	0.00(0.037)*	14
Contour	4	0.13(0.006)*	14	4	0.23(0.015)*	3	4	0.33(0.017)*	6	4	0.29(0.011)*	8
ENMC	8	0.17(0.006)*	8	5	0.22(0.014)*	4	9	0.34(0.016)*	5	9	0.39(0.012)*	3
Inter-part	5	0.14(0.007)*	10	5	0.22(0.017)*	5	5	0.36(0.015)*	4	**5**	**0.46(0.011)**	**1**
S	4	0.19(0.005)*	6	7	0.27(0.013)*	2	7	0.33(0.015)*	7	7	0.35(0.012)*	6
I	11	0.32(0.004)*	2	10	0.12(0.054)*	9	**10**	**0.57(0.009)**	**2**	9	0.18(0.018)*	11
S+I	**17**	**0.33(0.004)**	**1**	**12**	**0.27(0.017)**	**1**	**12**	**0.58(0.010)**	**1**	12	0.21(0.019)*	10
S+I(Ex)	43	0.24(0.008)*	5	1	0.10(0.033)*	11	26	0.37(0.025)*	3	30	0.42(0.019)*	2
Moments	1	0.16(0.004)*	9	1	0.02(0.017)*	12	1	0.33(0.012)*	8	1	0.03(0.037)*	13
CNN-A	57	0.29(0.010)*	3	56	0.17(0.027)*	7	10	0.03(0.036)*	14	18	0.26(0.045)*	9
CNN-G	18	0.14(0.009)*	11	18	0.10(0.022)*	10	16	0.08(0.018)*	13	11	0.38(0.023)*	4
CNN-F	252	0.25(0.010)*	4	79	0.14(0.026)*	8	16	0.10(0.017)*	12	44	0.07(0.139)*	12

Table 3: Performance of models on predicting average human accuracy in north/south categorisation, age, height and weight. All procedures were similar to Table 2 except that we trained regularised linear regression models instead of binary classifiers and model performance is the predicted vs observed correlations over test splits in from 10-fold cross validation, repeated 100 times. **R**-Rank, **Df**-degrees of freedom, **corr**-correlation between 10-fold cross validated predictions and obseved values for each rating.

References

[1] K. Brooks and O. Gwinn. No role for lightness in the perception of black and white? simultaneous contrast affects perceived skin tone, but not perceived race. *Perception*, 2010. 1, 2, 4

[2] X. Duan, C. Wang, Z. Liu, Xiang-dong Li, J. Wu, and H. Zhang. Ethnic features extraction and recognition of human faces. *2nd International Conference on Advanced Computer Control*, 2010. 1

[3] H. G. Feichtinger and T. Strohmer. *Gabor Analysis and Algorithms*. Birkhuser, 1998. 2

[4] S. Fu, H. He, and Z. G. Hou. Learning race from face: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(12):2483–2509, Dec 2014. 2

[5] E. Kocabey, M. Camurcu, F. Ofli, Y. Aytar, J. Marin, A. Torralba, and I. Weber. Face-to-bmi: Using computer vision to infer body mass index on social media. In *The International AAAI Conference on Web and Social Media (ICWSM)*, 2017. 6

[6] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Neural Information Processing Systems (NIPS)*, 2012. 2

[7] G. Levi and T. Hassner. Age and gender classification using convolutional neural networks. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) workshops*, June 2015. 4, 5

[8] S. Milborrow and F. Nicolls. Active shape models with sift descriptors and mars. *VISAPP*, 2014. 3, 4, 5

[9] T. Ojala, M. Pietikainen, and D. Harwood. Performance evaluation of texture measures with classification based on kullback discrimination of distributions. In *Proceedings of 12th International Conference on Pattern Recognition*, volume 1, pages 582–585 vol.1, Oct 1994. 2

[10] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *British Machine Vision Conference*, 2015. 4, 5, 6

[11] S. Setty, M. Husain, P. Beham, J. Gudavalli, M. Kandasamy, R. Vaddi, V. Hemadri, J. C. Karure, R. Raju, B. Rajan, V. Kumar, and C. V. Jawahar. Indian movie face database: A benchmark for face recognition under wide variations. In *2013 Fourth National Conference on Computer Vision, Pattern Recognition, Image Processing and Graphics (NCVPRIPG)*, pages 1–5, Dec 2013. 2

[12] R. Sharma and M. S. Patterh. Indian face age database: A database for face recognition with age variation. *International Journal of Computer Applications*, pages 21–28, 2015. 2

[13] G. Somanath, M. Rohith, and C. Kambhamettu. Vadana: A dense dataset for facial image analysis. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 2175–2182, Nov 2011. 2

[14] U. Tariq, Y. Hu, and T. S. Huang. Gender and ethnicity identification from silhouetted face profiles. In *2009 16th IEEE International Conference on Image Processing (ICIP)*, pages 2441–2444, Nov 2009. 1

[15] H. H. K. Tin and M. M. Sein. Race identification for face images. pages 118–120, Aug 2011. 1

[16] Y. Wang, H. Liao, Y. Feng, X. Xu, and L. Jiebo. Do they all look the same? deciphering chinese, japanese and koreans by fine-grained deep learning. *Arxiv*, 2016. 1, 2